GOODWIN

5 MAY 2026

Proceed With Caution: AI Platforms and Privilege

Rebecca Wardle, Alasdair McAlpine

Publicly available artificial intelligence (AI) tools are widely available and uptake in their use has grown exponentially in recent years. With use of such AI tools becoming increasingly normalised, it is important for individuals and businesses to be aware that by using them to conduct legal analysis, develop strategy, or draft in the context of a transaction, a commercial agreement, or dispute, they may lose their ability to assert privilege over them and — in a worst-case scenario — could risk having the exchanges disclosed in an used against them in future litigation or regulatory investigations.

In this alert, we summarise the principles of privilege in England, consider how the use of AI tools can trigger a loss of confidentiality and privilege under English law, and provide practical guidance for those seeking to use AI tools in a way that limits the risk of waiver.

Privilege Under English Law

Legal professional privilege is a fundamental right that allows a party to withhold certain documents and information from production to third parties, the court, regulators, and/or enforcement agencies. In summary, legal professional privilege encompasses:

- **Legal advice privilege**, which attaches to written or oral communications made (a) between a client and their lawyer (which includes in-house lawyers, provided they are acting in their capacities as lawyers); (b) under conditions of confidentiality; and (c) for the dominant purpose of enabling the client to seek, or the lawyer to give, legal advice.
- **Litigation privilege**, which attaches to written or oral communications made (a) between a client and their lawyer or between either of them and a third party; (b) under conditions of confidentiality; and (c) for the sole or dominant purpose of use in litigation that, at the time the communication is made, (i) is proceeding, pending, reasonably anticipated, or in contemplation and (ii) the client is, or reasonably anticipates becoming, a party to.

What Is the Effect of Uploading Privileged Information to AI Tools?

There is an important distinction between (i) publicly available and (ii) enterprise or contractually restricted AI tools, particularly with respect to data confidentiality and use. Most publicly available AI tools (including the free or consumer versions of widely used platforms) explicitly state that they retain and review user inputs for purposes such as improving the service, developing products, or training the underlying AI models.

Therefore, depending on the tool and its applicable terms, data uploaded to these platforms may be used for the benefit of the provider or otherwise be made accessible to third parties. In the absence of express contractual protections, it is prudent to assume that confidentiality will not be retained in any data uploaded to such tools. By contrast, enterprise or proprietary AI platforms typically operate under more restrictive data processing terms and practices. These tools generally provide that user data remains proprietary to the customer, will not be used to train the AI model for third-party use, and they usually offer additional controls around data retention, access, and security. Many AI providers offer both publicly available and enterprise versions of the same or similar tools, and the applicable data handling terms may differ significantly between them.

In a recent case, *UK and R (on the application of Munir) v. Secretary of State for the Home Department [2026] UKUT 81 (IAC)*, the Upper Tribunal (Immigration and Asylum Chamber) observed that:

> to put client letters and decision letters from the Home Office into an **open source AI tool** […] is to place this information on the internet in the **public domain**, and thus to **breach client confidentiality and waive legal privilege**. […] **Closed source AI tools** which do not place information in the public domain […] **are available for tasks** such as summarising **without these risks**. (Emphasis added)

In doing so, the Upper Tribunal essentially recognised that there is a distinction between AI platforms that retain and utilise data uploaded to them (typically the case for publicly available, free AI tools) and those that do not (typically enterprise or contractually restricted AI tools). It observed that uploading privileged documents to an unrestricted AI tool constitutes a breach of confidentiality and a loss of privilege. In comparison, uploading privileged documents to an enterprise or proprietary AI tool with contractual restrictions, which means uploaded information is not effectively placed in the public domain, does not lead to an immediate loss of confidentiality and privilege. Of course, we note that a loss of privilege could subsequently occur if confidentiality in the information was lost — for example, if the information was provided to or accessed by a third party.

The reasoning for this distinction is clear. Confidentiality is a prerequisite to both legal advice privilege and litigation privilege. Data uploaded to publicly available AI tools is typically uploaded on the basis that it may be used by the provider as previously discussed. As such, when that data is uploaded, it loses its confidentiality and, as a result, is no longer treated as privileged. Conversely, data uploaded to enterprise or proprietary platforms is generally (but this should always be checked) subject to contractual restrictions such that any uploaded data remains proprietary to the customer and is not used to train the AI model for third-party use. As such, this data retains its confidential and privileged status.

It is important to note that the Upper Tribunal in *Munir* was not seeing to dissuade parties from using AI tools. Rather, it recognised that — when used properly and responsibly — the tools available are “a step forward in legal practice.” Further, the Upper Tribunal’s observations on privilege do not amount to a binding ruling, not least because the substantive issues before it did not require a determination on the question of privilege. However, the comments are likely to prove influential because they consider the application of established principles of legal professional privilege in the context of the use of AI tools.

How Does This Compare to Guidance From Other Jurisdictions?

This is an evolving and fast-paced area of law, but other jurisdictions have reached similar decisions. In February 2026, a federal judge in the U.S. District Court for the Southern District of New York held that a criminal defendant's claims of privilege over his pre-indictment exchanges with a widely used AI chatbot were not privileged, paving the way for prosecutors to use that evidence against him in a fraud and embezzlement case. The defendant had discussed his intended legal strategy with the chatbot in advance of his indictment. The judge determined that these communications were not privileged, including because there is no reasonable expectation of privacy in AI communications. Further details in relation to this decision are available in our US alert.

Lessons for the Use of AI in a Legal Context

Clients may reasonably wish to upload advice from (for example) the company's external legal advisers to an AI tool to interrogate it further or to summarise it. However, the key takeaway is that clients must proceed with caution when uploading or copying and pasting confidential and legally privileged materials into AI tools. Unless an AI tool with sufficient privacy and data controls is used (typically an enterprise or proprietary tool), the uploaded information may lose confidentiality and, therefore, any privilege attached to it.

The arguments in relation to waiver of privilege arising from the use of unrestricted AI tools are potentially wide-ranging and apply irrespective of whether the individual using the tool is an in-house lawyer. If deemed to have been waived, then the client cannot then assert privilege over the materials that were uploaded to the tool. There is also a risk that the upload or exchange with the AI tool may trigger a collateral waiver over other related materials and/or documents, such that other documents and/or information that was not uploaded to the AI tool may subsequently be required to be disclosed to third parties, the court, regulators and/or enforcement agencies in the context of any future litigation or investigation concerning the same subject matter.

Dos and Don'ts

Do:

- Check the terms of service/use of any AI tool before inputting any confidential information or legal advice to ascertain how the information may be used by the AI tool.
- Adopt a cautious approach. If in any doubt about the platform and its terms, leave legally privileged and/or other confidential material out.
- Ask any law firms instructed what AI tools they use and what policies are in place to ensure there is no risk of a waiver of confidentiality.

Don't:

- Assume that an enterprise or proprietary AI tool is automatically safe to use. Review the data processing terms carefully or seek confirmation from the relevant person in the organisation that the AI tool has sufficient restrictions around the use of data uploaded to it.
- Assume that confidentiality or privilege is preserved because a document was marked "privileged and confidential" before it was uploaded to an AI tool — labelling makes no difference.

We would like to thank Hari Solomonides for their assistance with this alert.

CONTACTS

Rebecca Wardle
Partner
Financial Services
rwardle@goodwinlaw.com

Alasdair McAlpine
Associate
Investment Management
amcalpine@goodwinlaw.com